

PTAS for Ordered Instances of Resource Allocation Problems with Restrictions on Inclusions

Kamyar Khodamoradi^{*1}, Ramesh Krishnamurti^{†1}, Arash Rafiey^{‡1, 2} and Georgios Stamoulis^{§3}

¹School of Computing Science, Simon Fraser University, Burnaby, BC, Canada

²Department of Math and Computer Science, Indiana State University, Terre Haute, IN, USA

³Department of Data Science & Knowledge Engineering, Maastricht University, The Netherlands

kka50@sfu.ca, ramesh@sfu.ca
arash.rafiey@indstate.edu
georgios.stamoulis@maastrichtuniversity.nl

October 12, 2016

Abstract

We consider the problem of allocating a set I of m indivisible resources (items) to a set P of n customers (players) competing for the resources. Each resource $j \in I$ has a same value $v_j > 0$ for a subset of customers interested in j , and zero value for the remaining customers. The utility received by each customer is the sum of the values of the resources allocated to her. The goal is to find a feasible allocation of the resources to the interested customers such that for the Max-Min allocation problem (Min-Max allocation problem) the minimum of the utilities (maximum of the utilities) received by the customers is maximized (minimized). The Max-Min allocation problem is also known as the *Fair Allocation problem*, or the *Santa Claus problem*. The Min-Max allocation problem is the problem of Scheduling on Unrelated Parallel Machines, and is also known as the $R||C_{\max}$ problem.

In this paper, we are interested in instances of the problem that admit a Polynomial Time Approximation Scheme (PTAS). We show that an ordering property on the resources and the customers is important and paves the way for a PTAS. For the Max-Min allocation problem, we start with instances of the problem that can be viewed as a *convex bipartite graph*; a bipartite graph for which there exists an ordering of the resources such that each customer is interested in (has a positive evaluation for) a set of *consecutive* resources. We demonstrate a PTAS for the inclusion-free cases. This class of instances is equivalent to the class of bipartite permutation graphs. For the Min-Max allocation problem, we also obtain a PTAS for inclusion-free instances. These instances are not only of theoretical interest but also have practical applications.

1 Introduction and Problem Definition

In the general resource allocation problem, we are given a bipartite graph $H = (I, P, E)$ where the set of vertices I represents m indivisible resource items (or simply items) to be allocated to the set P of n players (or customers). The sets I and P form the two sets of the bipartition, and are indexed by numbers in $[m]$ and $[n]$ respectively. Thus, $I = \{x_1, x_2, \dots, x_m\}$ and $P = \{p_1, p_2, \dots, p_n\}$. Each player $p_j \in P$ has a utility

^{*}The author wishes to acknowledge support from a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant.

[†]The author wishes to acknowledge support from a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant.

[‡]The author is supported by COMPETE grant from Indiana State University

[§]The work of the author was partially supported by the Swiss National Foundation as part of the *Early Postdoc.Mobility* grant P1TIP2_152282

function $val_j(i) = v_i \geq 0$ for each item $x_i \in I$. This represents the value of item x_i for player p_j . If item x_i is adjacent to player p_j , then its value for her is $v_i > 0$ (v_i is a positive integer), otherwise its value is zero. This is represented in the graph H via the edge set E . If an item x_i has a value $v_i > 0$ to a player p_j , then there exists an edge $e \in E$ that connects x_i to p_j in H , and we say player p_j is adjacent to item x_i . If there is no edge between an item $x_i \in I$ and a player $p_j \in P$, then the item x_i has a value of zero to player p_j . Allocating an item $x_i \in I$ with value $v_i \in \mathbb{Q}_{>0}$, where $\mathbb{Q}_{>0}$ is the set of positive rational numbers, to a player adjacent to the item increases the utility of that player by v_i . Our goal is to find a feasible allocation of the resource items to the players that optimizes an objective function.

We assume there are no items of degree zero in H . If there exists such an item, we can safely remove it from the graph since it is not accessible to any player. In a feasible allocation (henceforth, *feasible assignment*) for n players, the set I is partitioned into n subsets, where each subset is allocated to the player with the same index. In other words, $I = I^1 \cup I^2 \cup \dots \cup I^n$, where items in I^j are allocated to player $p_j \in P$ who is adjacent to these items. Such a partition is denoted by $\mathcal{A} = (I^1, I^2, \dots, I^n)$.

The objective function that is optimized depends on the particular resource allocation problem we solve. In the Max-Min allocation problem, we seek a feasible solution that ensures that each player receives utility *at least* $t > 0$, assuming that we can guess the largest possible value of the parameter t . In other words, we seek a partition $\mathcal{A} = (I^1, \dots, I^n)$ such that $\sum_{i: x_i \in I^j} v_i \geq t, \forall j \in [n]$. A sufficiently good guess of t can be obtained by doing a binary search. It is worth mentioning that the assignments made in the Max-Min allocation problem do not necessarily need to be a partitioning of the items, meaning that one can leave some items unassigned. But since the goal is to maximize the sum of item values assigned, restricting the assignments to cover the entire set of items will not decrease the objective value. Therefore, to avoid further complications, we assume the assignments made are in fact partitions of the entire set of items. On the other hand, in the Min-Max allocation problem, we have item costs rather than item values. For instance, v_i 's can be the construed as processing times of a set of jobs and players can be processors, and we are required to assign every job to some processor. Naturally, we wish to allocate jobs to machines such that the overall finish time, or *make-span*, is minimized. Therefore, we seek a feasible solution that ensures each player receives a utility of *at most* $t > 0$ for the smallest t possible. In other words, we seek a partition $\mathcal{A} = (I^1, \dots, I^n)$ such that $\sum_{i: x_i \in I^j} v_i \leq t, \forall j \in [n]$.

In this paper, we study cases where we can obtain PTAS for these two resource allocation problems. We start with cases where the bipartite graph $H = (I, P, E)$ that models the problem is *convex*. A bipartite graph $H = (I, P, E)$ with two set of vertices I and P , is *convex* if there is an ordering $<_I$ of the vertices in I such that the neighbourhood of each vertex in P consists of consecutive vertices, *i.e.*, the neighbourhood of each vertex in P forms an interval. Indeed, we focus on *inclusion-free* instances which form a subset of the convex bipartite graphs. This subclass of graphs is also referred to as the class of *bipartite permutation graphs* and the class of *proper interval bigraphs*. We will discuss them in more detail in Section 2.1. Convex bipartite graphs (henceforward referred to as *convex graphs* for short) were introduced in [13] and are well known for their nice structures and both theoretical and practical properties. Many optimization problems become polynomial-time solvable or even linear-time solvable in convex graphs while remaining \mathcal{NP} -hard for general bipartite graphs [9]. Convex graphs can be recognized in linear time by using PQ-trees [9] and, moreover, the recognition algorithms provide the ordering $<_I$, given that the underlying graph is of course convex [8, 15, 23]. Various studies have considered the bipartite permutation graphs for other types of optimization problems. This class is of interest in graph homomorphism problems [14]. Also, there are recognition algorithms for bipartite permutation graphs [14, 29].

The interval case arises naturally in energy production applications where resources (energy) can be assigned and used within a number of successive time steps [18, 22]. That is, the energy produced at some time step is available only for a limited period corresponding to an interval of time steps. The goal is a fair allocation of the resources over time, *i.e.*, an allocation that maximizes the minimum accumulated resource we collect at each time step. In other words, we would like to have an allocation that guarantees the energy we collect at each time step is at least t , a pre-specified threshold. See also [28] for some applications in on-line scheduling.

1.1 Related Work

For the general Max-Min fair allocation problem, where a given item does not necessarily have the same value for each player (*i.e.*, every player has her own value for any item), no “good” approximation algorithm is known. In [7], by using the natural LP formulation for the problem and similar ideas as in [21], an additive ratio of $\max_{i,j} v_{ij}$ is obtained, which can be arbitrarily bad. A stronger LP formulation, the *configuration LP*, is used to obtain a solution at least opt/n in [5]. Subsequently, [4] provided a rounding scheme for this LP to obtain

a feasible solution of value no worse than $\mathcal{O}(\frac{\text{opt}}{\sqrt{n(\log^3 n)}})$. Recently in [26], an $\mathcal{O}(\sqrt{\frac{\log \log n}{n \log n}})$ approximation algorithm was proposed, which is close to the *integrality gap* of the configuration LP. On the negative side, there exists a simple $\frac{1}{2}$ inapproximability result [7] using the same ideas as in [21]. Due to the difficulty of the general case, subsequent research focused more on special cases. For the *restricted* case of the Max-Min allocation problem (also known as the restricted Santa Claus problem), where $v_{ij} \in \{0, v_j\}$ for $i \in [n]$ and $j \in [m]$, there is an $\mathcal{O}(\frac{\log \log \log n}{\log \log n})$ factor approximation algorithm [5]. Furthermore, the $\frac{1}{2}$ inapproximability result for the general case [7] carries over to this restricted case as well.

Recently, Feige proved that the integrality gap of the configuration LP defined and studied in [5], cannot be worse than some constant. In [3] an integrality gap of $\frac{1}{5}$ was shown for the same LP which was later improved to $\frac{1}{4}$. This implies that a $\frac{1}{4}$ -approximation algorithm based on rounding the configuration LP for the restricted Max-Min allocation is possible, although no such algorithm is available yet.

The authors provide a local search heuristic with an approximation guarantee of $\frac{1}{4}$. However, this heuristic was not known to run in polynomial time. Later, it was shown in [25] that the local search can be done in $n^{\mathcal{O}(\log n)}$ time. In [16] the authors provided a constructive version of Feige's original nonconstructive argument based on Lovász Local Lemma, thus providing the first constant factor approximation for the restricted Max-Min fair allocation problem. They provide an α -approximation algorithm for some *constant* α where an explicit value of α is not provided (but is thought to be a huge constant). Thus there is still a gap between the $\frac{1}{2}$ inapproximability result and the constant α approximability result in [16]. Very recently, a 13-approximation was given for the problem [2], which provides the first constant factor polynomial time approximation algorithm for the problem for a particular constant value. Their approach uses, in a highly non-trivial way, the local search technique for hypergraph matching that was used in [3]. Another important aspect of this approach is that the algorithm is purely combinatorial.

Several special cases of the Max-Min fair allocation problem have been studied. The case where $v_{ij} \in \{0, 1, \infty\}$ is shown to be hard in [19] and a trade-off between the running time and the approximation guarantee is established. In particular, for a number $\alpha \leq \frac{|I|}{2}$, it is shown how to design an $\alpha \cdot \frac{\text{opt}}{n}$ -approximation algorithm in time $|P|^{\mathcal{O}(1)}|I|^{\mathcal{O}(\alpha)}$. In [6] the authors consider the case in which each item has positive utility for a bounded number of players D , and prove that the problem is as hard as the general case for $D \leq 3$. They also provide a $\frac{1}{2}$ inapproximability result and a $\frac{1}{4}$ approximation algorithm for the *asymmetric* case (where each item has two distinct non-zero values for the two players interested in that item) when $D \leq 2$. The authors also provide a simpler LP formulation for the general problem and devise a polylogarithmic approximation algorithm that runs in quasipolynomial time (and a $|P|^\varepsilon$ -approximation algorithm that runs in $|P|^{\mathcal{O}(\frac{1}{\varepsilon})}$ time, for some $\varepsilon \geq 0$). The same result has been obtained in [10], which includes a $\frac{1}{2}$ approximation when $D \leq 2$, thus matching the bound proved in [6]. In [33], the author provides a PTAS for a (very) special case of the problem considered in this paper, namely, when the instance graph of the problem is a complete bipartite graph. In [22] a $\frac{1}{2}$ -approximation algorithm was developed for the class of instances considered in this paper, namely for the case where the intervals for each player are inclusion-free in the same sense introduced in this paper. See also [27], [24] for some other special cases that our results generalize.

The Min-Max-Allocation problem, or the $R||C_{\max}$ problem, as it is known in standard scheduling notation, is an important class of resource allocation problems. We seek an allocation (also known as an assignment) of jobs (the resources) to machines (the players) such that the makespan (the time by which the latest machine finishes its processing) is minimized.

For the $R||C_{\max}$ problem, a 2-approximation algorithm based on a characterization of the extreme point solutions of the natural linear programming relaxation of the problem is given [21]. The authors also provide a $\frac{3}{2}$ inapproximability result. These results have been adapted to the Max-Min case in [7]. So far, all efforts to improve either of the bounds have failed. In a very recent result [31], it is shown that the restricted version of $R||C_{\max}$ (where the processing time of a job j is v_j for a subset of the machines and infinity otherwise) admits an α approximation guarantee for α strictly less than 2. This result is an *estimation* result i.e., it estimates the (optimal) makespan of the restricted $R||C_{\max}$ within a factor of $\alpha = \frac{33}{17} + \varepsilon$ for some arbitrary small positive ε , although no polynomial time algorithm with such a performance guarantee is known. In [11], a 1.75 approximation algorithm is given for the restricted $R||C_{\max}$ problem where each job can be assigned to at most 2 machines. In this article, the authors claim that their 1.75 approximation for this restricted case is the first one that improves the factor 2 approximation of [21] on an unbounded number of machines and jobs. Our PTAS thus provides a certain strengthening of their claim, providing the first natural and non-trivial instance (as in the case of a complete bipartite graph) for which a PTAS is provided.

We note that further restrictions of this special case, where every job has a degree at most two, have been

studied [20]. Thus if the underlying bipartite graph is a tree, then a PTAS can be designed for the problem. If the processing times are in the set $\in \{1, 2\}$, then a $\frac{3}{2}$ -approximation algorithm exists, which is the best possible unless $\mathcal{P} = \mathcal{NP}$. Finally, [32] provides better approximations for several special cases of the problem. More importantly, it shows that the configuration LP for this restriction has an integrality gap of 2.

1.2 Our Contribution

We summarize our results below:

- (1) We start by presenting a PTAS for the restricted Max-Min fair allocation problem when the instance of the problem is an inclusion-free convex graph, that is, the neighborhood of each player is an interval of items (Section 2). Notice that this instance of the problem is (strongly) \mathcal{NP} -hard, as it contains complete bipartite graphs as a special case (each player is adjacent to all the items), which is known to be strongly \mathcal{NP} -hard [12].
- (2) Next, we modify our approach for the Max-Min allocation problem to obtain a PTAS for the $R || C_{\max}$ problem with inclusion-free convex graphs. (Section 3).

To obtain the PTAS for the instances considered in this paper, we first use scaling and rounding to classify the items into small and big items. Then, we prove that in a given instance of the problem, for any assignment of a certain value, another assignment of slightly less value but with simpler structure exists. Finally, we provide an algorithm that searches all these simple-structured assignments.

2 Max-Min Allocation Problem (Santa Claus) on Convex Graphs

An instance of the problem is given via a convex graph $H = (I, P, E)$ and a utility function $val : I \rightarrow \mathbb{Q}_{>0}$ which associates a value v_i to every item x_i in I . Structural properties and algorithms of convex graphs have received attention in the field of algorithmic graph theory. We refer the reader to texts such as [9] and [30]. For the sake of completeness, we give a definition of this class of graphs and introduce some notations that we will use throughout the paper.

2.1 Preliminaries and Notations

To define the class of convex graphs we first need to introduce the following property.

Definition 1 (Adjacency Property). [9] *Let $H = (X, Y, E)$ be a bipartite graph. An ordering $<_X$ of X in H has the adjacency property if for each vertex $y \in Y$, the neighbourhood of y in X denoted by $N(y)$ consists of vertices that are consecutive in $<_X$.*

A convex graph is defined as follows.

Definition 2 (Convex Graph). [9] *A bipartite graph $H = (X, Y, E)$ is convex if there exists an ordering of X or Y that satisfies the adjacency property.*

In this paper, we deal with a subclass of convex graphs known as inclusion-free convex graphs. In this subclass, inclusion may still occur between the intervals, but it should follow certain rules. We first explain the rules with the help of Figure 1. Assume that a convex graph $H = (I, P, E)$ is given as the input instance. Consider two players p and q in P . Their neighbourhoods in H can either be separate (their intersection is empty) or intersecting. If the neighbourhood of two players p and q is intersecting, four types of intersection between the two neighbourhoods may occur. These types are depicted in Figure 1. We say that two intervals are *properly overlapping* (Figure 1 (a)) if neither of the two intervals contains the other one, or *left inclusive* (Figure 1 (b)) if one interval contains the other and the two intervals share their left boundary, or *right inclusive* (Figure 1 (c)) if one interval contains the other and the two intervals share their right boundary, or *marginally-inclusive* (Figure 1 (d)) if one interval is completely contained in the other, and the intervals do not share their boundaries. The subclass of inclusion-free convex graphs forbids marginally-inclusion while left inclusion and right inclusion may still occur.

We define the class of inclusion-free convex graphs as follows. Given an ordering of the items $<_I$ that satisfies the adjacency property, for any pair of players if the neighbourhood of one player is completely included




Figure 1: Four types of intersecting intervals of two players

in the neighbourhood of the other, the smaller neighbourhood must either contain the leftmost or the rightmost item of the bigger neighbourhood with respect to the ordering $<_I$. This property is sometimes referred to as the *enclosure* property in the literature.

Definition 3 (Enclosure Property). *Let $H = (X, Y, E)$ be a bipartite graph. An ordering $<_X$ of X in H has the enclosure property if for each pair of vertices $y, y' \in Y$ for which $N(y) \subseteq N(y')$ in X , the vertices of $N(y') \setminus N(y)$ appear consecutively in $<_X$. We say that the graph H has the enclosure property if such an ordering of the vertices of H exists.*

As we mentioned earlier, every convex graph admits such an ordering and we can find it in linear time. Let $<_I$ be the ordering of items in set I that satisfies the adjacency property. For every vertex $p \in P$ let $[\ell_p, r_p]$ be the interval of the items adjacent to p . Based on the ordering $<_I$ on I , we define the following ordering on P .

Definition 4 (Lexicographical Ordering of Players). *For a given ordering $<_I$ of items I , an ordering of players is called lexicographical if and only if a player p is ordered before another player q whenever $\ell_p <_I \ell_q$, or $\ell_p = \ell_q$ and $r_p \leq_I r_q$ (breaking ties arbitrarily), in which $r_p \leq_I r_q$ means either $r_p = r_q$ or $r_p <_I r_q$.*

The adjacency property is in some contexts referred to as the *Min ordering* property. When a graph admits both the adjacency and the enclosure property, it is said to respect the *Min-Max ordering* property. In a given convex graph $H = (I, P, E)$ with the ordering of items $<_I$ that satisfies the adjacency property, let $<_P$ denote the lexicographical ordering of the players based on $<_I$. For arbitrary players $p, q \in P$ and arbitrary items $i, j \in I$, if p is adjacent to $i \in I$ and q is adjacent to j with $p <_P q$ and $j <_I i$, then by the adjacency property p is also adjacent to j , meaning that the neighbourhood of p contains the “minimum” of the two items i and j with respect to the ordering $<_I$. The enclosure property dictates that q also be adjacent to i , the maximum of the two items i and j , hence the name “Min-Max” ordering (not to be confused with Min-Max allocation problem) property. Note that not all convex graphs satisfy the enclosure (or Min-Max ordering) property. In fact, the class of convex graphs that satisfy the enclosure property is a proper subclass of the convex class and is known to be equivalent to two famous graph classes, the bipartite permutation graphs and the proper interval bigraphs.

The following observation shows an interesting property of graphs that satisfy both the adjacency and enclosure properties.

Observation 1. *Assume we are given an inclusion-free convex graph $H = (I, P, E)$ alongside an ordering of items $<_I$ that satisfies both the adjacency and the enclosure properties. If I is ordered according to $<_I$, then the lexicographical ordering of P based on $<_I$ also satisfies both properties. In other words, when P is ordered lexicographically based on $<_I$, (i) the neighbourhood of every item $x \in I$ forms an interval in P , and furthermore, (ii) for every pair of items $x_1, x_2 \in I$, their respective neighbourhoods are either properly overlapping, left inclusive, or right inclusive.*

Proof. We first show part (i). For the sake of contradiction, assume there is an item $x \in I$ such that the neighbourhood of x in P does not form an interval. This implies that a gap exists in the neighbourhood of x in P , meaning that there are players $p' <_P p <_P p''$ such that x is adjacent to p' and p'' but not p . Here, $<_P$ denotes the lexicographical ordering of players. Since the neighbourhood of players form intervals in I , the neighbourhood of p' can include x in one of the three ways depicted in Figure 2. The dotted edges represented the non-edges. We go through each case:

Case 1: x is in the middle of the interval of p' . As p is not adjacent to x , we first consider the case in which $\ell_p <_I x$. We can conclude $r_p <_I x$ since the neighbourhood of p is an interval in I . This means that




Figure 2: Different cases a gap may exist in the interval of an item x in an inclusion-free convex graph

either $\ell_p = \ell_{p'}$ which together with the fact that $r_p <_I r_{p'}$ (note that p' is adjacent to x and p is not, so the right boundary of the neighbourhood of p' must be to the right of that of p in this case) means that p should come before p' in the lexicographical ordering contradicting the assumption that $p' <_P p$, or we have that $\ell_{p'} <_I \ell_p$ and $r_p <_I r_{p'}$ which contradicts the assumption that the graph satisfies the enclosure property (the neighbourhood of p falls completely inside that of p'). Therefore it follows that $x <_I \ell_p$. In this case, since $p <_P p''$, we either have that $\ell_p = \ell_{p''}$ or $\ell_p <_I \ell_{p''}$. In both this situations, p'' cannot be adjacent to x , implying that the assumption of Case 1 over the neighbourhood of players is false.

Case 2: x is on the left boundary of the interval of p' . In this case, we certainly have that $x <_I \ell_p$ as p is lexicographically larger than p' and not adjacent to x . As in Case 1, either $\ell_p = \ell_{p''}$ or $\ell_p <_I \ell_{p''}$ due to the lexicographical ordering. In both situations, p'' cannot be adjacent to x which makes Case 2 a contradiction as well.

Case 3: x is on the right boundary of the interval of p' . Again, $x <_I \ell_p$, as assuming otherwise leads to a contradiction we mentioned in Case 1. Therefore, since $p <_P p''$, we get the two cases as before, $\ell_p = \ell_{p''}$ or $\ell_p <_I \ell_{p''}$, both leading to a contradiction.

The assumption that a gap exists in the neighbourhood of x leads to a contradiction in either case. This proves part (i) of the observation. For part (ii), we assume there is an item x whose neighbourhood is marginally included inside the neighbourhood of some other item x' . Again, there are two cases: one in which $x' <_I x$, and the other $x <_I x'$. We only show the former as the latter follows from symmetry. Since the interval of x is entirely inside the neighbourhood of x' without touching any of the boundaries, there must exist players $p' <_P p <_P p''$ such that p is adjacent to both x and x' , but p' and p'' are only adjacent to x' , lying on the left and right side of the neighbourhood of x . Figure 3 depicts the neighbourhood of players and items, with the dashed edges representing the non-edges. Note that the neighbourhood of each of the players is an interval in I . Players p' and p'' are both adjacent to x' but not x , while player p is adjacent to x . Since $x' <_I x$, we conclude that $\ell_{p'} <_I r_{p'} <_I r_p$ and $\ell_p <_I r_p <_I r_{p''}$. This contradicts the assumption that p'' is lexicographically greater than p . Therefore the intervals of items in the set P must be inclusion-free. \square

Given a convex graph $H = (I, P, E)$, and ordering $<_I$ of items that satisfies the adjacency property, and a target value t , our goal is to find a t -assignment of items to players, defined informally as a feasible assignment such that each player (machine) receives total utility at most t . More precisely,

Definition 5 (t -assignment for the Max-Min case). *A t -assignment, for any $t \geq 0$, is a feasible assignment such that every player p_j receives a set of items $I^j \subseteq [\ell_{p_j}, r_{p_j}]$ with total value at least t .*

In this paper, we only deal with the easier case of inclusion-free instances. All the definitions and theorems in later sections apply to the class of inclusion-free convex graphs (unless mentioned otherwise explicitly). To the best of our knowledge, no PTAS is known for general instances with margined-inclusive intervals, where the convex graph only satisfies the adjacency property. Whether there exists a PTAS for convex graphs, in general, is an interesting research question that we pose as an open problem.

2.1.1 A Note on the Notation

Throughout the paper, we use **bold** mathematical symbols to refer to vectors. Superscripts are reserved for player indices. For simplicity, we denote the induced subgraphs of the input convex graph H by $H' = (I', P')$




Figure 3: The contradictory assumption that intervals of items in the set of players P are not inclusion-free

for some $I' \subseteq I$ and $P' \subseteq P$ and refrain from explicitly mentioning the subset of edges $E' \subseteq E$. The set E' contains all those edges of E that are incident to at least one vertex in I' and at least one vertex in P' . The underline is used for emphasis. Finally, we use *italics* to highlight the terms we define in the paper.

2.2 Preprocessing the Instance

Given a particular instance of the problem, we first simplify the input instance. For a positive rational number t , we may assume that the value of each item is at most t . If item x_i has value $v_i > t$ then we set v_i to t without loss of generality. By a proper scaling, *i.e.*, dividing each value by t , we may assume that the value of each item is in $[0, 1]$. Observe that a t -assignment becomes a 1-assignment. We do a binary search to find the largest value of t for which a t -assignment exists (in the Min-Max case we seek the smallest t such that a t -assignment exists). The binary search is carried out in the interval $[0, \frac{1}{n} \sum_{i=1}^m v_i]$ where 0 is an absolute lower bound, and $\frac{1}{n} \sum_{i=1}^m v_i$ is an absolute upper bound of the optimal solution respectively (for the Min-Max case the binary search is carried out in the interval $[\frac{1}{n} \sum_{i=1}^m v_i, \sum_{i=1}^m v_i]$).

For a subset $P' \subseteq P$ of players, let $N_{[H]}(P')$ be the union of the set of all neighbours of the players in P' in the graph H . We let $N(P')$ denote this set whenever the graph H is implied by the context. For an interval $[i, j]$ of items, let $P[i, j]$ be the set of players whose entire neighbourhood lies in $[i, j]$, that is $P[i, j] = \{p \in P : N(p) = [\ell_p, r_p] \subseteq [i, j]\}$. For a subset $I' \subseteq I$ of items, let $val(I')$ denote the sum of values of all the items in I' . We note that in every 1-assignment, for every subset $P' \subseteq P$ of players, the sum of the value of items in its neighbourhood should be at least $|P'|$. In other words, $\forall P' \subseteq P : val(N(P')) \geq |P'|$. If the value of each item is 1, then this condition is the well known Hall's condition [17], a condition sufficient and necessary for a bipartite graph to have a perfect matching. We consider a more general version of the 1-assignment and derive a generalized version of Hall's condition below. Each player $p \in P$ has a demand $d(p)$. This version contains the 1-assignment as a special case, *i.e.*, the case where $d(p) = 1$ for all players $p \in P$. Also, for a subset of players $P' \subseteq P$, let $d(P')$ denote the sum of demands of all players in P' . Now the generalized Hall's condition (for the Max-Min case) becomes: $\forall P' \subseteq P : val(N(P')) \geq d(P')$. From now on we refer to this condition as Hall's condition. Note that this condition is necessary to satisfy the players' demands, but not sufficient (see Figure 4). Lemma 1 shows that in order to check Hall's condition for H it suffices to check it for every interval of items. Therefore, Hall's condition in our setting becomes Condition 2.1 below:

$$\forall [\ell, r] \subseteq [1, m] : \quad val([\ell, r]) \geq d(P[\ell, r]). \quad (2.1)$$

Lemma 1. *In order to check Hall's condition for H it suffices to verify Condition 2.1. In other words, it suffices to check Hall's condition for every set of players $P[\ell, r]$, $[\ell, r] \subseteq [1, m]$.*

Proof. Assume we are given a convex graph H . We claim that it is sufficient to check Hall's condition only for intervals of items, meaning that if there are any violations of Hall's condition, then there is at least one violation over an interval of items. In other words, we show that there exists a subset of the players P' for which (i) $N(P')$ consists of several maximal intervals (therefore $N(P')$ is not an interval itself) and (ii) Hall's condition is violated for P' if and only if there exists an interval of items $[\ell, r] \subseteq I$ for which Condition 2.1 is violated.

\Rightarrow : Assume that there exist several maximal intervals J_1, J_2, \dots, J_t whose union gives $N(P')$. Since each player is adjacent to an interval of items, there exists a corresponding partition of P' into subsets P'_1, P'_2, \dots, P'_t , such that $N(P'_i) = J_i$. Since Hall's condition is violated, we have

$$val(J_1) + \dots + val(J_t) = val(N(P')) < d(P') = d(P'_1) + \dots + d(P'_t)$$




Figure 4: An instance in which Hall's condition is satisfied for $t = 1$ but the optimal solution value is not greater than 0.4. In this examples, $d(p_1) = d(p_2) = d(p_3) = 1$.

Thus, there must exist an $i, 1 \leq i \leq t$, for which $\text{val}(J_i) = \text{val}(N(P'_i)) < d(P'_i)$. Now let ℓ and r be the leftmost and rightmost items in J_i respectively. Since $N(P'_i) = J_i = [\ell, r]$, then the neighbourhood of every player in P'_i should fall entirely in the interval $[\ell, r]$. Therefore, $P'_i \subseteq P([\ell, r])$. Thus, we conclude $\text{val}([\ell, r]) < d(P'_i) \leq d(P([\ell, r]))$, which is a violation of Condition 2.1 for the interval $[\ell, r]$.

\Leftarrow : Assume that Condition 2.1 is violated for an interval of items $[\ell, r] \subseteq [1, m]$, meaning that $\text{val}([\ell, r]) < d(P([\ell, r]))$. We have that $N(P([\ell, r])) \subseteq [\ell, r]$ since otherwise $P([\ell, r])$ contains a player whose entire neighbourhood does not fall in the set $[\ell, r]$, which contradicts the definition of $P([\ell, r])$. Let P' be $P([\ell, r])$. We then have that $\text{val}(N(P')) \leq \text{val}([\ell, r]) < d(P')$, which means that Hall's condition is violated for the set P' . This completes the second direction of the claim. \square

As a result of Lemma 1, it is sufficient to check Hall's condition for every interval of items. Since there are at most m^2 intervals, Hall's condition can be verified in time polynomial in the size of H .

2.2.1 Rounding the Instance

In a given instance H of the problem, we round the item values to obtain a rounded instance H' . For any integer $k \geq 4$, we let $\frac{1}{k}$ be the error parameter. For each instance for which there is an optimal 1-assignment, we seek an assignment such that each player receives a set of items with total value at least $1 - \frac{4}{k+1}$, $k \geq 4$.

- We call an item x_i *small* if its value is less than or equal to $\frac{1}{k}$. The values of small items are not changed in H' .
- We call an item *big* if its value is greater than $\frac{1}{k}$. If the value of item x_i , v_i , is in the interval $(\frac{1}{k}(1 + \frac{1}{k})^\tau, \frac{1}{k}(1 + \frac{1}{k})^{\tau+1}]$ for $\tau \geq 0$, it is rounded to $\frac{1}{k}(1 + \frac{1}{k})^{\tau+1}$ in H' .

After the rounding, there are at most $C = \lceil \log k / (\log(1 + \frac{1}{k})) \rceil$ categories of big items, that is items with distinct values more than $\frac{1}{k}$. One can easily verify that C is no more than $k^{1.4}$. For $0 \leq \tau \leq C$, let $q_\tau = \frac{1}{k}(1 + \frac{1}{k})^\tau$. For subset I' of I let $v_s(I')$ denote the value of the small items in I' . The rounding process in the Min-Max case is slightly different and we include it in the corresponding section.

2.3 The Algorithm for Inclusion-Free Convex Graphs

In this section, we consider the cases in which the neighbourhoods of no two players form a pair of margined-inclusive intervals. Note that left and right inclusion may still occur in such instances. The main result of this section is the following theorem.

Theorem 1. *Let H be an instance of the problem before rounding with n players and m items. Then for $k \geq 4$ there exists a $(1 - \frac{4}{k+1})$ -approximation algorithm for the Max-Min allocation problem on inclusion-free convex graphs with running time $\mathcal{O}((m+n)nm^{2(C+1)})$, in which $C \leq k^{1.4}$.*

Our proof technique builds upon a previous work of Alon et al. [1]. In particular, we use the notion of “input vectors” (Definition 6) in a dynamic programming algorithm that considers the players one by one and remembers plausible allocations for them that can lead to a $(1 - \varepsilon)$ -assignment. Input vectors are configuration vectors that indicate the number of available items in an instance of the problem or one of its subproblems. In their paper, Alon et al. study complete bipartite graphs, *i.e.*, the case where all items are available to all players. This fact allows them to ignore the actual assignment of items to players in their dynamic programming algorithm and only work with the configuration vectors for the set of available items, *i.e.*, the input vectors;

once the configuration vector for an assignment of items to a certain player is known, any set of items that matches the configuration vector can be assigned to the player. In other words, all items of a certain size are identical to the algorithm. Since the number of all possible input vectors is polynomial in the instance size, they manage to provide a polynomial-time algorithm in the end.

Our work focuses on a more general case, that of inclusion-free convex graphs, which contains complete bipartite graphs as a subset. In this setting, an item may be adjacent only to a subset of players. As a result, to identify a subproblem, we cannot simply represent the set of available items to each player with the corresponding input vector. On the other hand, one cannot consider all possible ways of assigning items to players since the number of such assignments is exponential and computationally prohibitive. For the case of margined inclusion-free instances, we prove one can still identify the subproblems from their respective input vectors. To that end, we introduce $(1 - \varepsilon)$ -assignments of a certain structure (which we call simple-structured assignments) and prove their existence in any instance of the problem which admits a 1-assignment. Such assignments are defined in Definitions 7 and 11, and shown to exist in Lemmas 2 and 3.

In this section, we assume that the convex graph H is a rounded input. For a rounded instance, it suffices to set the constant parameter ε to $\frac{3}{k}$. In the proof of Theorem 1, we show how this bound guarantees a $1 - \frac{4}{k+1}$ approximation factor for the original instance. The dynamic programming algorithm has a forward phase and a backward phase. In the forward phase, the algorithm checks the existence of a solution by filling up a table. During this phase, the algorithm runs n steps, one for each player. At a given step j , it checks for feasible $(1 - \varepsilon)$ -assignments $\mathcal{A}' = (I^{n-j+1}, I^{n-j+2}, \dots, I^n)$ for players $p_{n-j+1}, p_{n-j+2}, \dots, p_n$, called *partial* assignments, based on the solutions of the sub-problems in the previous step (feasible $(1 - \varepsilon)$ -assignments for players p_{n-j+2}, \dots, p_n). During the first step ($j = 1$), the algorithm searches for $(1 - \varepsilon)$ -assignments for player p_n from scratch, therefore this case is treated differently. The backward phase also has n steps. During this phase, the algorithm considers the players in the reverse order of the forward phase, *i.e.*, from p_1 to p_n , and generates the final set of items to be allocated to each player based on the information stored in the table.

Every partial assignment of the forward phase indicates a potential assignment of items to players*. After every such assignment, the set of available items changes, thus a subproblem is instantiated. The induced subgraph of H on the set of available items and unsatisfied players is called a *remainder* graph. The algorithm regards every remainder graph as an instance of a subproblem of the original allocation problem. Note that the number of possible partial $(1 - \varepsilon)$ -assignments and their corresponding remainder graphs can be exponential. Some remainder graphs may eventually result in a solution and some may not. Thus, at each step, a naïve dynamic programming approach might need to remember the remainder graphs, which will lead to an exponential memory (and time) complexity. By exploiting the properties of convex graphs (specifically the Min-Max ordering) we can show that one can do better than the naïve approach. In this section, we examine some structural properties of the problem that allow us to find a solution for any given instance while checking a polynomial[†] number of simple-structured assignments. These properties will also be useful for a more general case of the problem studied in Section ?? . Our dynamic programming algorithm uses input vectors, which we introduce here.

Definition 6 (Input Vector). *For a given convex graph H for an input instance with ν_τ big items of size q_τ for $\tau = 1, 2, \dots, C$ and small items of total value in the interval $(\frac{\nu_0-1}{k}, \frac{\nu_0}{k}]$, an input vector is a configuration vector of the form $\nu(H) = \boldsymbol{\nu} = (\nu_0, \nu_1, \dots, \nu_C)$.*

Note that for an arbitrary input vector $\boldsymbol{\nu} = (\nu_0, \nu_1, \dots, \nu_C)$, $\nu_\tau \leq m$ for all big items of size q_τ , $1 \leq \tau \leq C$ as there can be at most m items of a certain value. Furthermore, ν_0 is the total value of the small items in integral multiples of $\frac{1}{k}$ which has been rounded up. Therefore, $\nu_0 \leq m$ since $\frac{m}{k}$ is an upper bound on the total value of the small items. As a result, there are at most m^{C+1} possible input vector values, which is a polynomial in the size of the problem instance. In the rest of this paper, we let \mathcal{V} denote the set of all possible input vectors. Similar to the algorithm by Alon et al., we are interested in a dynamic programming approach that only deals with the input vectors in \mathcal{V} rather than the (potentially exponential) remainder graphs. Unlike the case of complete graphs, we cannot assign items of a certain size interchangeably. However, by imposing the right set of restrictions on the assignments we can retrieve (reconstruct) a remainder graph from any input vector whose total sum of items is almost equal to the sum of values indicated by the input vector. In particular, we wish to show the following two facts:

Fact 2.1. *Whenever there exists an arbitrary 1-assignment of items to players for an instance of the problem, there also exists a restricted $(1 - \varepsilon)$ -assignment.*

*The assignment will be finalized once the algorithm finishes the steps of the forward phase.

†In the number of players and items, but exponential in the inverse of the error parameter.

Fact 2.2. *Given that a 1-assignment exists for an instance of the problem, then there also exists a polynomial-time algorithm that, given the input vectors for every step (partial assignment) of a $(1 - \varepsilon)$ -assignment, reconstructs the remainder graph of every step, only from its input vector and not knowing the actual partial assignment, in such a way that the total value of the items in each reconstructed remainder graph is*

- *the same as that of the corresponding remainder graph in the original (and unknown) 1-assignment for every category $\tau = 1, 2, \dots, C$ of the big items and,*
- *only a small fraction ε more than its counterpart in the original remainder graph for the small items.*

As the first of the restrictions we impose on the assignment to obtain remainder graphs of simpler structures, we define *right-aligned* assignments.

Definition 7 (Right-Aligned Assignment). *For a given convex graph H , an ordering of items $<_I$ that satisfies the adjacency property, and a lexicographical ordering of players with respect to $<_I$, an assignment of items to players $\mathcal{A} = (I^1, I^2, \dots, I^n)$ is called right-aligned if and only if it satisfies the following properties recursively:*

- (1) *Let p_n be the last (rightmost) player in the graph. All the big items of value $q_\tau = \frac{1}{k}(1 + \frac{1}{k})^\tau$ ($1 \leq \tau \leq C$) assigned to p_n are the rightmost ones with that value in $N_{[H]}(p_n)$. Furthermore, all the small items assigned to p_n are also the rightmost ones in her neighbourhood.*
- (2) *In the remainder graph $H' = (I \setminus I^n, P \setminus p_n)$, that is if p_n and all the items assigned to her are removed from the graph H , the rest of the assignment, $\mathcal{A}' = (I^1, I^2, \dots, I^{n-1})$, is a right-aligned assignment.*
- (3) *An empty assignment (i.e., when no players are left in the graph) is considered a right-aligned assignment.*

Next, we prove Fact 2.1 for the right-aligned assignments in a lemma called the *Alignment Lemma* (Lemma 2). Intuitively, the lemma states that if a 1-assignment exists, then there exists a $(1 - \varepsilon)$ -assignment with items aligned to the right, in which $\varepsilon = \frac{1}{k}$. We first need the following observation and definition.

Observation 2. *Convex graphs have the hereditary property, meaning that every induced subgraph of a convex graph is also convex.*

Proof. Assume otherwise, i.e., there exists no ordering of items that satisfies the adjacency property for a subgraph $H' = (I', P')$ of a convex graph $H = (I, P, E)$, where $P' \subseteq P$ and $I' \subseteq I$. Since H is convex, there exists an ordering $<_I$ for I which satisfies the adjacency property. On the other hand, $<_I$, when induced on the subset of items I' does not satisfy the adjacency property in H' . Let $<_{I'}$ denote this induced ordering. There must exist a player $p \in P'$ and items $x, y, z \in I'$ for which x and z are in the neighbourhood of p in H' and y is not while $x <_{I'} y <_{I'} z$. If we consider these items in H , x and z are still in the neighbourhood of p and y is still not, while $x <_I y <_I z$. This contradicts the assumption that $<_I$ satisfies the adjacency property in H . \square

Definition 8 (Assignment Vector). *For an assignment of items to players $\mathcal{A} = (I^1, I^2, \dots, I^n)$, and for a player p_j , $1 \leq j < n$, let $\alpha^j = \nu(H^j)$ for $j = 1, 2, \dots, n$, in which $H^j = (P \setminus \{p_n, p_{n-1}, \dots, p_{j+1}\}, I \setminus (I^n \cup I^{n-1} \cup \dots \cup I^{j+1}))$ is the remainder graph for player p_j in the assignment \mathcal{A} (for $j = n$, let $H^n = H$). We call the vector $\alpha^j = (\alpha_0^j, \alpha_1^j, \dots, \alpha_C^j)$ the assignment vector of player p_j in the assignment \mathcal{A} . Furthermore, let $\alpha = (\alpha^1, \alpha^2, \dots, \alpha^n)$ be the assignment vector for the entire assignment \mathcal{A} .*

Lemma 2 (The Alignment Lemma). *Suppose there exists a 1-assignment for a given convex graph H . Then, there also exists a minimal $(1 - \frac{1}{k})$ -right-aligned assignment for H . Furthermore, the two assignments will have identical assignment vector.*

Proof. We let $\mathcal{A} = (I^1, \dots, I^n)$ denote the 1-assignment stated in the lemma and let α be its assignment vector. Also, we denote by I_τ^j the set of big items of value q_τ assigned to player p_j for $1 \leq \tau \leq C$ and by I_0^j , the set of small items allocated to p_j in the assignment \mathcal{A} . We prove the lemma by transforming the assignment \mathcal{A} into a minimal $(1 - \frac{1}{k})$ -right-aligned assignment, $\mathcal{A}' = (J^1, J^2, \dots, J^n)$. \mathcal{A}' is minimal in the sense that the removal of any items from the sets J^1, \dots, J^n will cause its value to drop below $1 - \frac{1}{k}$. This task is carried out in two independent and consecutive rounds. In round 1, the big items assigned to each player are aligned to the right while the small items are aligned in round 2. Each round consists of n steps, one for each player. The two rounds are independent of each other in that the alignment of the big items in the first round has no impact on the alignment of small items in the second round. In the following, we assume the players are ordered lexicographically based on the ordering of items $<_I$.

Round 1: In round 1, the big items are right-aligned in n steps, where in each step, the big items of a single player are aligned in C micro-steps, one micro-step for each item category. We initialize j to be n , indicating that we start with the last player, p_n , and let τ be 1. Also, let H^j denote the remainder graph at the beginning of step j . At the τ^{th} micro-step, we only look at items of size q_τ in the graph and replace the set I_τ^j with its right-aligned counterpart J_τ^j . Note that according to Observation 2, the induced subgraph on items of a certain size, say H_τ^j , forms a convex graph. We start with $J_\tau^j = I_\tau^j$. If I_τ^j is right-aligned, *i.e.*, all items of value q_τ assigned to p_j are the rightmost ones in $N_{H_\tau^j}(p_j)$, then we return J_τ^j . If however, I_τ^j is not right-aligned, there must be two items x_r and x_t of the same value q_τ , where $x_r <_I x_t$ and x_r is assigned to p_j , but x_t is not. If x_t is not assigned to any other player, then we can simply assign it to p_j instead of x_r . If it is assigned to some player p_ℓ , x_r must also be connected to p_ℓ . Otherwise p_ℓ would come after p_j in the lexicographical ordering of the players. Since the two items have the same value, we use the adjacency property to swap them in the assignment. Thus, p_j gets x_t and p_ℓ gets x_r in the transformed assignment J_τ^j . Now we have one less item out of alignment. We continue this process until there are no items in J_τ^j out of alignment. Then we proceed to the next micro-step by setting $\tau \leftarrow \tau + 1$. At the end of C micro-steps, we let J^j be $\bigcup_{\tau=1}^C J_\tau^j$. Furthermore, we obtain the remainder graph $H^{j-1} = (I \setminus J^j, P \setminus \{p_j\})$. This graph will be the input to the next step of the alignment. We now proceed to the next step by updating $j \leftarrow j - 1$. Note that the assignment vector for $\mathcal{A}' = (J^1, J^2, \dots, J^n)$, say α' , is identical to α on the big item coordinates since for any arbitrary player p , we did not change the number of big items of any category assigned to p . As a result, we have not lost any solution quality.

Round 2: In round 2, we obtain a right-aligned assignment for small items. We first explain the procedure and then prove its correctness. Let H_0 denote the subgraph of H that contains only small items. The neighbourhood of each player in H_0 is still an interval (by Observation 2). For $j \in [n]$, let $1 - w_j$ be the total value of the big items assigned to player p_j in the 1-assignment. Therefore, w_j denotes the deficit of player p_j that should be satisfied by small items. If the demands of each player p_j is set to w_j for $j \in [n]$, then H_0 must satisfy Hall's condition with these demands. This is because a 1-assignment exists in the original graph for which Hall's condition is necessary, and that during the first round, the values of the big items assigned to the players did not change. Following the idea of round 1, we align the items in n steps. We start with the last player, so we let j be n and let I_0^j denote the set of small items assigned to p_j . We also let H_0^j denote the remainder graph at the beginning of step j . Note that $val(I_0^j) = w_j$. We replace this set with the set of right-aligned items, J_0^j , of value at least $w_j - \frac{1}{k}$. We first let J_0^j be the empty set. Then, we pick small items from the neighbourhood of p_j in H_0^j and add them to J_0^j . We continue to do this as long as $val(J_0^j) \leq w_j - \frac{1}{k}$. When the algorithm stops, since the value of each small item is at most $\frac{1}{k}$, $w_j - \frac{1}{k} < val(J_0^j) \leq w_j$. By doing this, we ensure that player p_j gets a value strictly greater than $w_j - \frac{1}{k}$ in small items, and strictly greater than $1 - \frac{1}{k}$ in total. The remainder graph for the next step of round 2, H_0^{j-1} , is obtained by $H_0^{j-1} = (I \setminus J_0^j, P \setminus \{p_j\})$. Note that each such remainder graph is still a convex graph because we started with a convex graph H_0 and at each step, we removed items from right to left in the ordering $<_I$. Also, note that the small coordinate of the assignment vectors of \mathcal{A} and \mathcal{A}' (α and α' respectively) are identical too, which means $\alpha = \alpha'$. The reason is that we packed the small items in \mathcal{A}' in such a way that they have the same total value as in \mathcal{A} when counted in integral multiples of $\frac{1}{k}$ and rounded up. We then move to the next step by updating $j \leftarrow j - 1$.

To prove the correctness, we show that after each step j , $1 \leq j \leq n$, we can obtain a right-aligned assignment of small items the sum of whose values is strictly greater than $w_j - \frac{1}{k}$ for player p_j . Recall that we chose every J_0^j ($1 \leq j \leq n$) to be the minimal right-aligned set of items assigned to player p_j for which $w_j - \frac{1}{k} < val(J_0^j) \leq w_j$. For the sake of contradiction, assume that for some j , $1 \leq j \leq n$, we cannot provide a right-aligned assignment of small items J_0^j for which $val(J_0^j) > w_j - \frac{1}{k}$. Let t be the largest such index. This assumption implies that $val(N_{[H_0^t]}(p_t)) \leq w_t - \frac{1}{k}$. We let $H_0^{t,n}$ denote the induced subgraph of H_0 on players p_t, p_{t+1}, \dots, p_n and the small items in their neighbourhoods. We consider the partial right-aligned assignments of small items in $H_0^{t,n}$ represented by $J_0^t, J_0^{t+1}, \dots, J_0^n$, and claim that for a subset of the players p_t, p_{t+1}, \dots, p_n , Hall's condition is violated for small items in the original graph H_0 . Since Hall's condition is necessary for any feasible assignment that fulfills the demands of players, this in turn implies that there is no 1-assignment for the instance, which contradicts our earlier assumption. To prove this claim, we introduce the notion of a *gap* in the assignment. With respect to an ordering of items $<_I$, a gap exists in a (partial) assignment if x_j is not assigned to any player, but there exists another item x_i , such that $x_i < x_j$ in $<_I$ and x_i is assigned to some player. The item x_j is said to be in the gap with respect to the partial assignment. Based on this notion, we consider two cases:

Case 1: *There exists a gap in the partial assignment.* Let u be the smallest index in $N_{[H_0^{t,n}]}(p_t, p_{t+1}, \dots, p_n)$

for which item x_u is in the gap for the partial right-aligned assignment. By the choice of u , x_{u-1} must be assigned to a player in the set $\{p_t, p_{t+1}, \dots, p_n\}$. Let p_v be that player. Note that $x_u \notin N_{[H_0^{t,n}]}(p_v)$ as otherwise x_u would be assigned to p_v in the right-aligned assignment. Item x_{u-1} is thus the rightmost item assigned to player p_v . Let x_ℓ and x_r be the left-most and the right-most items in $N_{[H_0^{t,n}]}(p_t, p_{t+1}, \dots, p_v)$ respectively (note that $r = u - 1$). Then,

$$\begin{aligned}
val([\ell, r]) &= \sum_{\substack{j \in [\ell, r] : \\ x_j \in H_0^{t,n}}} val(x_j) \\
&= \sum_{j=t}^v val(J_0^j) \\
&\leq val(J_0^t) + \left(\sum_{j=t+1}^v w_j \right) \\
&\leq w_t - \frac{1}{k} + \left(\sum_{j=t+1}^v w_j \right) \\
&< \sum_{j=t}^v d(j).
\end{aligned}$$

where the first inequality is due to the fact that the total value assigned to each player p_j in the right-aligned assignment is at most w_j , the second inequality is due to the choice of p_t , and last inequality holds since $d_j = w_j$ for all players p_j . Thus, Hall's condition is not satisfied for the set of players p_t, p_{t+1}, \dots, p_v in \hat{H} .

Case 2: *There are no gaps in the partial assignment.* Every item in $N_{[H_0^{t,n}]}(p_t, p_{t+1}, \dots, p_n)$ has been assigned in the partial assignment. Thus $\bigcup_{j=t}^n J_0^j = N_{[H_0^{t,n}]}(p_t, p_{t+1}, \dots, p_n)$. Once again, let x_ℓ and x_r be the left-most and the right-most items in $N_{[H_0^{t,n}]}(p_t, p_{t+1}, \dots, p_v)$ respectively. Similarly,

$$\begin{aligned}
val([\ell, r]) &= \sum_{j \in [\ell, r]} val(x_j) \\
&\leq v(J_0^t) + \left(\sum_{j=t+1}^n w_j \right) \\
&\leq \left(\sum_{j=t}^n w_j \right) - \frac{1}{k} \\
&< \sum_{j=t}^n d(j).
\end{aligned}$$

This implies that Hall's condition is violated in H_0 for the set of players p_t, p_{t+1}, \dots, p_n .

□

Unfortunately, Fact 2.2 is not true for the right-aligned assignment restriction, as shown in the following example.

Example 2.3. *In the example shown in Figure 5, each of the circle items is considered small items and have a value of $\frac{1}{10}$ and each of the square items, which are called big items, have a value of $\frac{1}{4}$. Now, consider two different right-aligned 1-assignments.*

- (1) *In the first one, player p_3 gets the 5 rightmost circle items in her neighbourhood as well as the 2 rightmost square items, player p_2 also gets the 5 rightmost circle items in her neighbourhood alongside the 2 rightmost*

Proof. This is because if in any right-aligned assignment, we take the unassigned private items in the remainder graphs and allocate them to their respective owners, we can only increase the value received by each player. Therefore, Fact 2.1 holds for non-wasteful right-aligned assignments as well. \square

Lemma 3 proves Fact 2.2 for this type of restricted assignments.

Lemma 3. *Assume there exists a non-wasteful right-aligned 1-assignment $\mathcal{A} = (I^1, I^2, \dots, I^n)$ for an instance of the problem. Also assume we are given the original inclusion-free convex graph H and the input vectors for partial assignments $\mathcal{A}^j = (I^j, I^{j+1}, \dots, I^n)$ for $j = n-1, n-2, \dots, 1$, but not the partial assignments themselves. There exists a polynomial time algorithm that can reconstruct the remainder graphs for every \mathcal{A}^j in such a way that the total value of items in each reconstructed graph is at most $\varepsilon = \frac{2}{k}$ more than the total value of the items in the actual corresponding remainder graph for \mathcal{A}^j .*

Proof. We prove the lemma by providing the reconstruction algorithm. The algorithm is, in fact, a simple left-to-right sweep[‡] of the items. We prove the correctness in Claim 1. For $j = n, n-1, \dots, 1$, we generate a remainder graph after each partial assignment $\mathcal{A}^j = (I^j, I^{j+1}, \dots, I^n)$, using only the input vectors. Let $\nu = (\nu_0, \nu_1, \dots, \nu_C)$ be the input vector of the remainder graph after this partial assignment. Let $\hat{H} = (\hat{I}, \hat{P})$ denote the remainder graph that our algorithm outputs at the end of this step. We set \hat{P} to be $\{p_1, p_2, \dots, p_j\}$. As for the set of items \hat{I} , for the big items of value q_τ , $\tau = 1, \dots, C$, we select the left-most ν_τ items in the graph H and add them to \hat{I} . For small items, we select a maximal set of left-most items whose total value is less than $\frac{\nu_0+1}{k}$. This method of selection ensures that the items that have been assigned to the players thus far are right-aligned. Furthermore, if at any point, stranded items (of degree 0) appear in the remainder graph, we will return NULL. This ensures that private items are not left unassigned. Note that since \mathcal{A} is assumed to be a non-wasteful right-aligned assignment, so are all its partial assignments as well. Next, we show the correctness of the algorithm by proving Claim 1.

Claim 1. *In a given inclusion-free instance H , for a non-wasteful right-aligned partial assignment $\mathcal{A}^j = (I^j, I^{j+1}, \dots, I^n)$, the set of big items in \hat{H} , the remainder graph reconstructed by the right-to-left sweep algorithm, is identical to that of the original remainder graph of \mathcal{A}^j , and the set of small items of \hat{H} is a superset of the set of small items of the original remainder graph. Furthermore, the sum of values for small items in \hat{H} is at most $\frac{2}{k}$ more than that of the original remainder graph.*

Proof. We treat big items and small items separately.

Big items: First, we consider the induced graph on the set of big items of a certain size q_τ for some arbitrary $\tau = 1, \dots, C$. Assume that $\hat{P} = \{p_1, p_2, \dots, p_j\}$ for some $1 \leq j < n$ (the case where $j = n$ is trivial since the remainder graph would be the original graph H). Also let H' denote the original remainder graph for \mathcal{A}^j . For contradiction, assume that the graphs H' and \hat{H} are not identical when induced on the set of big items of value q_τ . Since both graph H' and \hat{H} have the same number of items of value q_τ , there must exist items x' and \hat{x} such that x' appears in H' , but not in \hat{H} and \hat{x} belongs to \hat{H} , but not to H' . This means that the item \hat{x} is assigned to a player p_x , $x \geq j+1$ in one of the earlier partial assignments $\mathcal{A}^{j+1}, \mathcal{A}^{j+2}, \dots, \mathcal{A}^n$. The item x' is either adjacent to a player p_y , $1 \leq y \leq j$, or it is an item of degree 0 in H' and \hat{H} (note that players p_{j+1}, \dots, p_n do not belong to the graphs H' and \hat{H}). If the latter case occurs, it means that x' is a stranded item in H' , contradicting the fact that H' is the remainder graph of a non-wasteful right-aligned assignment. In fact, when faced with a stranded item, the reconstruction algorithm returns NULL which signals the dynamic programming algorithm to skip the current input vector as it does not correspond to a valid remainder graph. Therefore, we can assume there exists a player p_y adjacent to x' . Since $y \leq j$ and $x \geq j+1$, we have that $p_y < p_x$. By the adjacency property, we have that $\hat{x} \in N_{[H]}(p_y)$. p_x is not adjacent to x' since otherwise the item \hat{x} assigned to her would be out of alignment. The item x' is assigned to the player p_y , therefore we conclude that $r_{p_x} < r_{p_y}$. Since $p_y < p_x$, it must be the case that $\ell_{p_y} < \ell_{p_x}$. This is a contradiction to the assumption that there are no inclusions in H . This situation is depicted in Figure 6.

Small items: Assume that the set of small items of \hat{H} is not a superset of the set of small items of H' , meaning that an item x' exists such that x' belongs to H' but not to \hat{H} . If there exists another item \hat{x} such that \hat{x} belongs to \hat{H} , but not to H' , then, similar to the situation in Figure 6, we conclude that a margined inclusion should exist in H , which is a contradiction. Therefore, we may assume that \hat{H} is a subgraph of H' . Recall that $\nu = \nu(H') = (\nu_0, \nu_1, \dots, \nu_C)$ and the sum of values of all the small items in H' is at most $\frac{\nu_0}{k}$.

[‡]Note that the items assigned are aligned to the right, and naturally, the remaining items in the graph would be aligned to the left.




Figure 6: If the set of big items of H' and \hat{H} are not identical, a margined inclusion should exist. The arrow indicates assignment in one of the predecessors of H' .

(including x'). The total value of small items in \hat{H} is at least $\frac{\nu_0}{k}$ by definition. This is a contradiction since we assumed \hat{H} is a subgraph of H' . Therefore, the set of small items of \hat{H} must be a superset of H' . In the proposed graph \hat{H} , we selected the minimal set of small items whose value is greater than or equal to $\frac{\nu_0}{k}$. Such a set may have a value as high as $\frac{\nu_0+1}{k} - \varepsilon$ for some small value of $\varepsilon > 0$ since the value of small items is less than or equal to $\frac{1}{k}$. The minimum total value of small items in H' is no less than $\frac{\nu_0-1}{k}$ by the definition of input vectors, hence the difference of at most $\frac{2}{k}$. This complete the proof of the claim. \square

\square

Observation 4. *Note that, as mentioned in Claim 1, the algorithm tends to overestimate the value of small items still available in the graph by at most an additive constant factor of $\frac{2}{k}$, but it never underestimates. Therefore when a player p_j expects a bundle of small items worth d_j , she might receive $d_j - \frac{2}{k}$ worth of small items instead.*

Based on Observation 4, we next provide a dynamic programming algorithm that enumerates all non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignments for a rounded instance H . From Lemma 2 and Observation 3 one can conclude that whenever a 1-assignment exists for a given instance of the problem, so does a non-wasteful right-aligned $(1 - \frac{1}{k})$ -assignment. Therefore, if we only consider non-wasteful right-aligned assignments instead of all assignments, we are still able to find an assignment with objective value at least $(1 - \frac{1}{k})$, provided that a 1-assignment for the instance exists in the first place. Lemma 3 and Observation 4 together show that for any non-wasteful right-aligned assignment for an instance of the problem, an approximation of the assignment can be retrieved, and we only lose an additive fraction of $\frac{2}{k}$ in solution quality. Therefore, the general idea of the algorithm is to consider all the m^{C+1} possible input vectors in a dynamic programming algorithm, retrieve a non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignment for each vector, and mark the feasible one. If it eventually manages to find at least one feasible assignment for p_1 (the last player considered by the algorithm), then it reports success. Before presenting the algorithm, we first introduce two functions that are invoked by the algorithm:

- (1) **retrieve:** Function `retrieve` takes as input parameters input vector ν and player index j , and returns remainder graph $H' = (I', P')$. The graph $H' = (I', P')$ that it reconstructs is such that $P' = (p_1, p_2, \dots, p_j)$ and I' is formed by the left-to-right sweep algorithm in Lemma 3[§]. Function `retrieve` returns `NULL` if at any point during reconstruction, there is at least one stranded item in the remainder graph.
- (2) **feasible:** Function `feasible` takes as input parameters remainder graph H' and set of items I' , and returns either true or false. It returns true if each of the following is true: H' is not `NULL`, set I' is entirely in the neighbourhood of the last player in H' , and the sum of item values in I' is at least $1 - \frac{3}{k}$. Otherwise, Function `feasible` returns false.

Algorithm Assignment (for Inclusion-Free Convex Instances): Algorithm Assignment uses dynamic programming to fill the entries of an $n \times m^{C+1}$ table, in which each row represents a player, and each column represents an input vector (we assume an arbitrary ordering of all valid input vectors). Let the matrix M denote this table. Each entry (j, ν) of M is composed of two fields:

[§]In fact, Function `retrieve` also needs the input inclusion-free convex graph H as the third input parameter. To simplify notation, we omit this parameter from the function calls.

- (1) $M(j, \nu).bit$ is a binary variable. If $M(j, \nu).bit$ is set to 1, then the input vector ν is called *marked* for player p_j . Otherwise, the vector ν is *unmarked* for player p_j .
- (2) $M(j, \nu).ptr$ either contains NULL or a reference to another input vector ν' (in the beginning, this pointer is initialized to NULL for every entry).

Let $\nu_{in} = (\nu_0, \nu_1, \dots, \nu_C)$ be the input vector for the original graph H . The algorithm starts the forward phase with the last player p_n . For every input graph ν (i.e., every column of the table M), it forms the remainder graph $\hat{H} = \text{retrieve}(\nu, n-1)$. Let \hat{I} denote the set of items in \hat{H} and also let $I^n = I \setminus \hat{I}$. The set I^n provides a potential assignment for player p_n . The algorithm checks if this assignment is a non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignment for p_n using the boolean function $\text{feasible}(H, I^n)$. Note that for the first step of the algorithm, the original graph H serves as the remainder graph. We call input vector ν a *successful* input vector for remainder graph H' and player p_n if p_n is the last player in H' and a call to the function $\text{feasible}(H', I^n)$ returns **true**. All the successful assignments are marked by setting $M(n, \nu).bit = 1$. For a successful input vector ν for player p_n , we set $M(n, \nu).ptr = \nu_{in}$, while for unsuccessful vectors, the value of $M(n, \nu).ptr$ is not changed. In this section, we will not use the field **ptr** in the forward phase of the algorithm[¶].

After this first iteration for player p_n , the algorithm proceeds to a similar procedure for other players. For a player p_j , $1 \leq j \leq n-1$, it seeks $(1 - \frac{3}{k})$ -assignments once more. The only difference is that since the assignment should provide a $(1 - \frac{3}{k})$ -assignment for all players p_n, p_{n-1}, \dots, p_j , for every entry (j, ν) the algorithm looks at all marked entries of the previous row, $j+1$. Let $M(j+1, \nu')$ be such an entry. The algorithm retrieves the graphs $\hat{H} = \text{retrieve}(\nu, j-1)$ and $\hat{H}' = \text{retrieve}(\nu', j)$. We let \hat{I} and \hat{I}' denote the item sets of graph \hat{H} and \hat{H}' respectively. If the set of items $I^j = \hat{I}' \setminus \hat{I}$ is confirmed to be at least a $(1 - \frac{3}{k})$ -assignment for player p_j by function $\text{feasible}(\hat{H}', I^j)$, the entry $M(j, \nu).bit$ is set to 1. Furthermore, if $M(j, \nu).ptr$ is NULL, it is updated to ν' to mark this step of the assignment, indicating that the input vector ν was achieved from a remainder graph represented by the input vector ν' . If the function feasible returns **false**, the algorithm simply moves to the next marked entry of the previous row. This procedure continues until the algorithm either finds a successful input vector for p_j , or that there is no successful vector in the previous row. In the former case, it reports success and moves to the backward phase in which the actual assignment is retrieved. In the latter case, it simply reports failure. In the backward phase, the non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignment can be obtained by following the **ptr** back from the last row to the first. Before formally presenting the algorithm, we make the following remarks.

Remark 1. If Function *feasible* returns **false** at any stage, it may be due to any one of the following three reasons. Either (i) the set I^j does not have enough total item value, or (ii) there are items in set I^j which are not in the neighbourhood of player p_j , or (iii) the remainder graph \hat{H} returned by *retrieve* contains stranded items. In each of these cases, the input vector ν does not represent a valid vector according to the set of restrictions imposed on the assignments. Therefore, the algorithm should ignore ν and simply move to the next input vector.

Remark 2. If the **ptr** field is already assigned a value, we do not change it since for the purpose of finding a feasible $(1 - \frac{1}{3})$ -assignments for margined-inclusive free instances, it is not important how we arrived at an input vector ν in the current iteration. If a vector ν represents an assignment that satisfies players p_n, \dots, p_{j+1} , a simple retrieval procedure that removes items from the right (as suggested in Lemma 3) can retrieve a remainder graph that is a close enough approximation of the true remainder graph.

Remark 3. In selecting the potential assignment I^j for player p_j at any iteration, we may assign a value more than the target value $1 - \frac{3}{k}$. Because we do not wish to leave items stranded or unaligned, the algorithm allocates every item in the difference of two consecutive remainder graphs to the player for whom the iteration is run. Nonetheless, as we formally discuss in Lemma 4, this is not at odds with our purpose of guaranteeing bundles of value $1 - \frac{3}{k}$ or higher in the rounded instance for every player.

The pseudocode for the algorithm is provided below. In the following, a vector ν is said to be less than or equal to ν' if every entry of ν is less than or equal to its counterpart in ν' .

In Lemma 4 next, we show that Algorithm Assignment finds a $(1 - \frac{3}{K})$ -assignment in a rounded instance if one exists. We use Observation 5 in its proof.

Observation 5. For a given instance H , assume H' and H'' are two remainder graphs, and that H' is an induced subgraph of H'' . Let $\hat{H}' = \text{retrieve}(\nu(H'), j')$ and $\hat{H}'' = \text{retrieve}(\nu(H''), j'')$, where j' and j'' are the indices of the last players in H' and H'' respectively. Then, \hat{H}' is also a subgraph of \hat{H}'' .

[¶]We will however make use of this field in the forward phase of the algorithm in Section ??

Algorithm 1 Algorithm Assignment for Inclusion-Free Instances (Forward Phase)

Data: rounded instance H .

Result: returns either a $(1 - \frac{3}{k})$ -assignment for all players in the rounded instance, or report failure.

```
for every vector  $\nu$  in  $\mathcal{V}$  such that  $\nu \leq \nu_{in}$  do
     $\hat{H} \leftarrow \text{retrieve}(\nu, n-1)$ 
     $\hat{I} \leftarrow$  the set of items in  $\hat{H}$ ;  $I^n \leftarrow I \setminus \hat{I}$ 
    if feasible( $H, I^n$ ) then
         $M(n, \nu).bit \leftarrow 1$ 
        if  $M(n, \nu).ptr == \text{NULL}$  then
             $M(n, \nu).ptr \leftarrow \nu_{in}$ 
        end
    end
end
end
for  $j \leftarrow n-1$  downto 1 do
    for every vector  $\nu$  in  $\mathcal{V}$  such that  $\nu \leq \nu_{in}$  do
        for every vector  $\nu'$  in  $\mathcal{V}$  such that  $\nu' \geq \nu$  and  $M(j+1, \nu').bit == 1$  do
             $\hat{H} \leftarrow \text{retrieve}(\nu, j-1)$ ;  $\hat{H}' \leftarrow \text{retrieve}(\nu', j)$ 
             $\hat{I} \leftarrow$  the set of items in  $\hat{H}$ ;  $\hat{I}' \leftarrow$  the set of items in  $\hat{H}'$ ;  $I^j \leftarrow \hat{I}' \setminus \hat{I}$ 
            if feasible( $\hat{H}', I^j$ ) then
                 $M(j, \nu).bit \leftarrow 1$ 
                if  $M(j, \nu).ptr == \text{NULL}$  then
                     $M(j, \nu).ptr \leftarrow \nu'$ 
                end
            end
        end
    end
end
end
```

Proof. The function `retrieve` preserves the set of players and the set of big items. Therefore, to prove the claim of the observation, it is enough to show that the set of small items in \hat{H}' is a subset of the set of small items in \hat{H}'' . Let $\nu' = \nu(H')$ be $(\nu'_0, \nu'_1, \dots, \nu'_C)$ and $\nu'' = \nu(H'')$ be $(\nu''_0, \nu''_1, \dots, \nu''_C)$. The fact that H' is a subgraph of H'' implies that $\nu'_0 \leq \nu''_0$, which means that the function `retrieve` has not packed more small items from the left to \hat{H}' compared to \hat{H}'' . Therefore, the set of small items of \hat{H}' is a subset of that of \hat{H}'' . \square

Lemma 4. Let H be a rounded instance of the problem with n players and m items. The Algorithm Assignment assigns (in polynomial time) to each player a set of items with value at least $1 - \frac{3}{k}$ if a 1-assignment exists for all players. The algorithm returns failure if no $(1 - \frac{1}{k})$ -assignment exists for H .

Proof. From Lemma 2, we know that where there exists a 1-assignment for a rounded instance H , there is also a right-aligned $(1 - \frac{1}{k})$ -assignment for the instance, which in turn implies the existence of a non-wasteful right-aligned $(1 - \frac{1}{k})$ -assignment. Now, assume the algorithm forms a bundle of items I^j to be potentially assigned to an arbitrary player p_j . The way the algorithm forms the bundles is by finding the difference in the item sets of two remainder graphs, one before the assignment is made and one after. In doing so, it guesses the remainder graphs from their corresponding input vectors. Further, assume that the two remainder graphs before and after assignment are H' and H'' respectively, thus H'' is a subgraph of H' . Also, the remainder graphs guessed by the function `retrieve` are \hat{H}' and \hat{H}'' . Based on Observation 5, \hat{H}'' is also a subgraph of \hat{H}' . In a worst case scenario, the function `retrieve` may guess the total value of small items in \hat{H}'' to be at most $\frac{2}{k}$ more than the actual value (the value of small items in H''), while it guesses the same quantity correctly for \hat{H}' (that is, equal to the total sum of small item values in H'). This means that the difference set I^j would have $\frac{2}{k}$ less in small item values than the right-aligned $(1 - \frac{1}{k})$ -assignment implied by the original remainder graphs H' and H'' . Thus, the set I^j results in a $(1 - \frac{3}{k})$ -assignment for player p_j . Therefore, for every possible non-wasteful right-aligned $(1 - \frac{1}{k})$ -assignment to players, the algorithm considers a non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignment instead. Since the existence of the former results in the existence of the latter, Algorithm Assignment is guaranteed to find a $(1 - \frac{3}{k})$ -assignment to players if a $(1 - \frac{1}{k})$ -assignment exists. On the other hand, when the algorithm reports failure, we can be certain that a $(1 - \frac{1}{k})$ -assignment does not exist, since

otherwise, our algorithm would be able to find a slightly worse non-wasteful right-aligned $(1 - \frac{3}{k})$ -assignment that is guaranteed to exist. \square

Now we can prove Theorem 1.

Proof of Theorem 1. Each player receives a set of items with value at least $1 - \frac{3}{k}$ for the rounded instance (see Lemma 4). Each item value is rounded up by at most a multiplicative factor of $1 + \frac{1}{k}$. Therefore each player receives a set of items with value at least $(1 - \frac{3}{k})/(1 + \frac{1}{k}) = 1 - \frac{4}{k+1}$ of the optimal. The two inner **for** loops of Algorithm Assignment take $m^{2(C+1)}$. The outer loop is run for n players, and finding the remainder graphs and potential bundles need time $\mathcal{O}(m + n)$. Thus, the Algorithm Assignment runs in time $\mathcal{O}((m + n)nm^{2(C+1)})$. \square

3 Min-Max Allocation Problem $(R \mid \mid C_{\max})$ On Convex Graphs

In this section, we will show that an adaptation of the techniques used in the previous sections, will give us a PTAS for the Min-Max fair allocation of jobs to machines in a Min-Max allocation problem. Most of the techniques used in this section are very similar to those of Section 2, therefore we treat the proofs more concisely here, only highlighting the modifications and changes.

3.1 Problem Definition and Preliminaries

As in instance of this problem, we are given a set $M = \{M_1, M_2, \dots, M_m\}$ of identical machines or processors and a set $J = \{J_1, J_2, \dots, J_n\}$ of jobs. Each job J_j has the same processing time p_j on a subset of machines and it has processing time ∞ on the rest of the machines. The goal is to find an assignment of the (entire set of) jobs to the machines $\mathcal{A} = (J^1, J^2, \dots, J^m)$, such that the maximum load among all the machines is minimized. Formally, we are given a convex graph $H = (M, J, E)$ where M is a set of machines and J is a set of jobs and E denotes the edge set. There is an edge in E between a machine and a job if the job can be executed on that machine. Also given is a utility function $p : J \rightarrow \mathbb{Q}_{>0}$ which assigns a processing time to each job. With slight abuse of notation, we write $p(J_j)$ as p_j for short, and job J_j requires p_j units of processing time to run to completion on any machine M_i , provided that there exists an edge between J_j and M_i in the set E . If the edge does not exist, the processing time of J_j on M_i is not bounded.

Similarly, as before, we assume that the input bipartite graph H satisfies the adjacency and the enclosure properties, *i.e.*, is an inclusion-free convex graph. In other words, we assume that we have an ordering $<_J$ of jobs (which orders the jobs as J_1, J_2, \dots, J_n) such that each machine can execute consecutive jobs (an interval of jobs). More formally, we denote the interval of job J_i by $[\ell_i, r_i]$. We assume that J_i is before J_j and write $J_i <_J J_j$ whenever $\ell_i < \ell_j$ or $\ell_i = \ell_j, r_i \leq r_j$. Based on $<_J$, an ordering of jobs that satisfies the adjacency property, we define a lexicographical ordering of the machines in the same manner as Section 2. Also as before, we may assume that $0 \leq p_i \leq 1$ by scaling down the processing time values.

Remark 4. *Similar to the case of Max-Min allocations, if the input convex graph H is inclusion-free, meaning that it satisfies the enclosure property as well as the adjacency property, then not only the neighbourhood of every machine is an interval of jobs, but also every job is adjacent to a consecutive set of machines. More precisely, given the ordering of jobs $<_J$ that satisfies the adjacency and enclosure property and assuming that the set of machines M is ordered lexicographically based on $<_J$, then the neighbourhood of each job $J_j \in J$ forms an interval in the set M . This can be seen as a consequence of Observation 1.*

Hall's Condition: Hall's condition needs to be slightly modified for the case of Min-Max allocations. For subset J' of jobs let $w(J')$ and $w_s(J')$ be the sum of the processing times of all jobs and small^{||} jobs in J' respectively. Also let $N_{[H]}(J')$ (or $N(J')$ when the graph H' is implied by context) denote the neighbourhood of the subset of jobs J' in M . A necessary condition for having a maximum load which is at most 1 is that for every subset J' of machines $w(J') \leq |N(J')|$. More generally, assume each machine $M_i, 1 \leq i \leq m$, has a maximum *allowable* load denoted by $a(M_i)$. Also, for a subset of machines M' , let $a(M')$ be the sum of allowable loads of all the machines in M' . Then, Hall's condition for the Min-Max allocation problem states that in order to have an assignment in which every machine has a load below its allowable maximum, it is necessary to have $w(J') \leq a(N(J'))$ for every subset of jobs J' . For the case of Max-Min allocations on inclusion-free graphs, we only need to check this condition for every interval of machines.

^{||}We use definitions for *small* and *big* jobs similar to the definitions for *small* and *big* items in Section 2.2.1.

Lemma 5. *In order to check Hall's condition for Min-Max allocations in an inclusion-free graph H , it is enough to check the following condition:*

$$\forall [\ell, r] \subseteq [1, m]: w(\mathcal{J}([\ell, r])) \leq \sum_{i=\ell}^r a(M_i) \quad (3.1)$$

in which $\mathcal{J}([\ell, r])$ denotes the set of jobs whose entire neighbourhood of machines falls in the interval $[\ell, r]$. We refer to Condition 3.1 as Hall's condition for Min-Max allocations.

Proof. The proof mostly follows the proof of Lemma 1 except for the fact that it uses the property mentioned in Remark 4. Assume we are given an inclusion-free convex graph H . As before, we show that there exists a subset of the jobs J' for which (i) $N(J')$ is not an interval and therefore can be represented as the union of several maximal intervals, and (ii) Hall's condition is violated if and only if there exists an interval of machines $[\ell, r]$ for which Condition 3.1 is violated.

\Rightarrow : Assume that $w(J') > a(N(J'))$ and that there exist several maximal intervals of machines $\mathcal{N}_1, \mathcal{N}_2, \dots, \mathcal{N}_t$ whose union gives $N(J')$. Since each job is adjacent to an interval of machines by Remark 4, there exists a corresponding partition of J' into subsets J'_1, J'_2, \dots, J'_t , such that $N(J'_i) = \mathcal{N}_i$. Since Hall's condition is violated, we have

$$a(N_1) + \dots + a(N_t) = a(N(J')) < w(J') = w(J'_1) + \dots + w(J'_t)$$

Thus, there must exist an $i, 1 \leq i \leq t$, for which $a(\mathcal{N}_i) = a(N(J'_i)) < w(J'_i)$. Let ℓ and r be the leftmost and rightmost machines in $N(J'_i)$. We have that $a(N(J'_i)) = \sum_{i=\ell}^r a(M_i) < w(J'_i) \leq w(\mathcal{J}([\ell, r]))$ since $w(J'_i) \subseteq w(\mathcal{J}([\ell, r]))$.

\Leftarrow : Assume that Condition 3.1 is violated for an interval of machines $[\ell, r] \subseteq [1, m]$, meaning that $w(\mathcal{J}([\ell, r])) > \sum_{i=\ell}^r a(M_i)$. Let J' be $\mathcal{J}([\ell, r])$. Then, $N(J') \subseteq [\ell, r]$ since otherwise it means there exists a job in $J' = \mathcal{J}([\ell, r])$ whose neighbourhood stretches beyond the interval $[\ell, r]$, contradicting the definition of $\mathcal{J}([\ell, r])$. Therefore $w(J') > \sum_{i=\ell}^r a(M_i) \geq a(N(J'))$. Therefore, Hall's condition is violated for the set J' . \square

3.2 Rounding the Instance

Before running the algorithm on the instance, we round the processing times of jobs based on an input error parameter k . First of all, by properly scaling the values we assume every job has a processing time less than or equal to one. We say a job is *small* if its processing time is less than or equal to $\frac{1}{k}$ after scaling, and we say a job is *big* if its processing time is strictly larger than $\frac{1}{k}$. As opposed to the Max-Min case, we round down the processing times for big jobs for the Min-Max case. Thus, if p_j , the processing time of job J_j , is in the interval $[\frac{1}{k}(1 + \frac{1}{k})^i, \frac{1}{k}(1 + \frac{1}{k})^{i+1})$, then it is replaced by $\frac{1}{k}(1 + \frac{1}{k})^i$. Using this method, we obtain at most $C = \lceil \frac{\log k}{\log(1 + \frac{1}{k})} \rceil$ distinct processing times for big jobs (each of which is more than $\frac{1}{k}$, the maximum processing time for small jobs). For $1 \leq \tau \leq C$, let $q_\tau = \frac{1}{k}(1 + \frac{1}{k})^i$. Each q_τ , $1 \leq \tau \leq C$, denotes the rounded processing time of category τ of big jobs.

3.2.1 The Algorithm for Inclusion-Free Convex Graphs

The main theorem of this section is the following.

Theorem 2. *Let H be an instance of the problem before rounding with n jobs and m machines. Then, for $k \geq 4$ there exists a $(1 + \frac{4}{k} + \frac{3}{k^2})$ -approximation algorithm for the Min-Max allocation problem on inclusion-free convex graphs with the running time of $\mathcal{O}((m+n)mn^{2(C+1)})$ in which $C \leq k^{1.4}$.*

To show this result, we borrow from the techniques used in Section 2 extensively. We use the notions of *t-assignment*, *input vector*, *assignment vector*, and *right-aligned* assignments in the same sense as we did for the Max-Min allocations (jobs are allocated to machines here, instead of items to players). We modify the definition of input vectors slightly for Min-Max allocations below.

Definition 12 (*t-assignment for the Min-Max case*). *A t-assignment, for any $t \geq 0$, is a feasible assignment such that every machine M_i receives a set of jobs $J^i \subseteq [\ell_{M_i}, r_{M_i}]$ with total processing time at most t .*

Definition 13 (*Input Vector*). *For a given convex graph H for an input instance with ν_τ big jobs of processing time q_τ for $\tau = 1, 2, \dots, C$ and small jobs of total processing time in the interval $[\frac{\nu_0}{k}, \frac{\nu_0+1}{k})$, an input vector is a configuration vector of the form $\nu(H) = \boldsymbol{\nu} = (\nu_0, \nu_1, \dots, \nu_C)$.*

Then, we can prove the following *alignment* lemma.

Lemma 6 (The Alignment Lemma for Min-Max Allocations). *Suppose there exists a 1-assignment for a given convex graph H . Then, there also exists a maximal $(1 + \frac{1}{k})$ -right-aligned assignment for H . Furthermore, the two assignments will have an identical assignment vector.*

Proof. We assign the jobs in two rounds and make use of the adjacency property. During the first round, the big jobs are aligned to the right in the same manner as Lemma 2. In doing so, we do not change the assignment vector since the same number of big jobs as the 1-assignment are allocated to the machines. During the second round, we assign the small items. The only difference between the assignment of jobs and the assignment of items is that we would prefer to allocate as many jobs as possible to a machine as long as the makespan is below the threshold of $1 + \frac{1}{k}$. Therefore, we pack maximal sets of small jobs and assign them to the machines. Due to the way the assignment vectors are defined, we still have not changed the vector after this phase. \square

By the problem definition, we are not allowed to have stranded jobs in the Min-Max allocation problem. Therefore, every right-aligned assignment that we make is also a *non-wasteful* right-aligned assignment. The following two facts hold for any non-wasteful right-aligned assignment of jobs to machines.

Fact 3.1. *Whenever there exists an arbitrary 1-assignment of jobs to machines for an instance of the problem, there also exists a restricted $(1 + \varepsilon)$ -assignment, in the sense that all the jobs assigned are restricted to be aligned to the right.*

Fact 3.2. *Given that a 1-assignment exists for an instance of the problem, there also exists a polynomial-time algorithm that, given the input vectors for every step (partial assignment) of a $(1 + \varepsilon)$ -assignment, reconstructs the remainder graph of every step in such a way that the total processing time of the jobs in each reconstructed remainder graphs is*

- *the same as that of the corresponding remainder graph in the original (and unknown) 1-assignment for every category $\tau = 1, 2, \dots, C$ of big jobs and,*
- *only a small fraction ε less than its counterpart in the original remainder graph for small jobs.*

While Lemma 6 verifies Fact 3.1 for non-wasteful right-aligned assignments, the following lemma ensures that Fact 3.2 also holds for these types of assignments.

Lemma 7. *Assume there exists a non-wasteful right-aligned 1-assignment $\mathcal{A} = (J^1, J^2, \dots, J^m)$ for an instance of the problem. Also assume we are given the original inclusion-free convex graph H and the input vectors for partial assignments $\mathcal{A}^j = (J^j, J^{j+1}, \dots, J^m)$ for $j = m - 1, m - 2, \dots, 1$, but not the partial assignments themselves. There exists a polynomial time algorithm that can reconstruct the remainder graphs for every \mathcal{A}^j in such a way that the total processing time of jobs in each reconstructed graph is at most $\varepsilon = \frac{2}{k}$ less than the total processing time of jobs in the actual corresponding remainder graph for \mathcal{A}^j .*

Proof. The proof uses a left-to-right sweep again. We do not provide the details here since it is identical to the proof of Lemma 3, except for the way small jobs are chosen in the reconstruction process. For small jobs, we select a minimal set of left-most jobs whose total processing time is greater than $\frac{\nu_0 - 1}{k}$ and add them to the remainder graph. Since the processing time of each small job is less than $\frac{1}{k}$, this ensures that the sum of processing times of the jobs we choose in the reconstructed remainder graph is less than or equal to that of the original remainder graph, but the deficit is not more than $\frac{1}{k}$. The rest of the proof follows the proof of Lemma 3. \square

Algorithm Assignment (for the Min-Max Problem on Inclusion-Free Convex Instances): Algorithm Assignment for the Min-Max allocation problem is a dynamic programming algorithm similar to the algorithm given in Section 2. Again we make use of two utility functions, namely the reconstruction function *retrieve* and the boolean function *feasible*. Function *retrieve* reconstructs a remainder graph based on an input vector and a player index which potentially has slightly less total processing time than the original remainder graph of a partial assignment, and function *feasible* checks whether a given assignment of jobs is feasible in a remainder graph and does not leave any stranded jobs behind. Using these functions, the algorithm fills in the entries of a $m \times n^{C+1}$ table T . Each entry (i, ν) of the table has two fields:

- (1) $T(i, \nu).bit$: this binary value is set to 1 if the input vector ν is marked for machine i , and set to 0 otherwise.
- (2) $T(i, \nu).ptr$: contains either NULL or a reference to another input vector ν' in the previous row, intended for use in the backward phase of the algorithm where an actual assignment of the jobs to the machines is retrieved from the table T .

The algorithm is as follows.

Algorithm 2 Algorithm Assignment for Inclusion-Free Instances (Forward Phase)

Data: rounded instance H .

Result: returns either a $(1 + \frac{3}{k})$ -assignment for all machines in the rounded instance, or report failure.

```

for every vector  $\nu$  in  $\mathcal{V}$  such that  $\nu \leq \nu_{in}$  do
     $\hat{H} \leftarrow \text{retrieve}(\nu, m - 1)$ 
     $\hat{J} \leftarrow$  the set of jobs in  $\hat{H}$ ;  $J^m \leftarrow J \setminus \hat{J}$ 
    if  $\text{feasible}(H, J^m)$  then
         $T(m, \nu).bit \leftarrow 1$ 
        if  $T(m, \nu).ptr == \text{NULL}$  then
             $T(m, \nu).ptr \leftarrow \nu_{in}$ 
        end
    end
end
for  $i \leftarrow m - 1$  downto 1 do
    for every vector  $\nu$  in  $\mathcal{V}$  such that  $\nu \leq \nu_{in}$  do
        for every vector  $\nu'$  in  $\mathcal{V}$  such that  $\nu' \geq \nu$  and  $T(i + 1, \nu').bit == 1$  do
             $\hat{H} \leftarrow \text{retrieve}(\nu, i - 1)$ ;  $\hat{H}' \leftarrow \text{retrieve}(\nu', i)$ 
             $\hat{J} \leftarrow$  the set of jobs in  $\hat{H}$  ;  $\hat{J}' \leftarrow$  the set of jobs in  $\hat{H}'$ ;  $J^i \leftarrow \hat{J}' \setminus \hat{J}$ 
            if  $\text{feasible}(\hat{H}', J^i)$  then
                 $T(i, \nu).bit \leftarrow 1$ 
                if  $T(i, \nu).ptr == \text{NULL}$  then
                     $T(i, \nu).ptr \leftarrow \nu'$ 
                end
            end
        end
    end
end

```

Lemma 8. Let H be a rounded instance of the problem with n jobs and m machines. Algorithm Assignment assigns (in polynomial time) to each machine a set of jobs with value at most $1 + \frac{3}{k}$ if a 1-assignment exists for all machines. The algorithm returns failure if no $(1 + \frac{1}{k})$ -assignment exists for H .

Proof. By Lemma 6, a $(1 + \frac{1}{k})$ -assignment exists whenever there is a 1-assignment for the problem instance. Using the retrieve function, we may guess the remainder graph of any of the partial assignments to have a deficiency in total processing time which is no more than $\frac{2}{k}$ when compared to the original remainder graph. This together with the error caused by the alignment imply the total sum of processing times of the jobs assigned to each machine is at most $\frac{3}{k}$ more than the same value in the 1-assignment (if a successful assignment is identified). It is also straightforward to see that whenever there exists no $(1 + \frac{1}{k})$ -assignment for the instance, the algorithm fails to recover a feasible assignment for a machine, resulting in a failure report. This completes the correctness proof of the Algorithm Assignment. \square

In the last part, we prove Theorem 2 using the lemmas in this section.

Proof of Theorem 2. First we discuss the approximation guarantee. Through alignment and reconstruction of the remainder graphs, we may pack an extra $\frac{3}{k}$ units of processing time in the bundle of jobs assigned to a machine. This extra value is magnified by a multiplicative factor of $1 + \frac{1}{k}$ in the rounded instance. Therefore, the approximation guarantee would be $(1 + \frac{3}{k}) \cdot (1 + \frac{1}{k}) = 1 + \frac{4}{k} + \frac{3}{k^2}$. We now discuss the running time of

the algorithm. The two inner loops of the Algorithm Assignment take a combined total time of $n^{2(C+1)}$ while the outer loop runs for each machine for a total of m times. Retrieving the remainder graph can be done in $\mathcal{O}(m+n)$, hence the running time is $\mathcal{O}((m+n)n^{2(C+1)})$ in total. The correctness is already shown in Lemma 8. \square

4 Conclusion and Future Work

In all instances of the problem considered in this paper, a proper ordering has played an important role. We conjecture that a dynamic programming algorithm similar to the one used in this paper provides a PTAS for the case of margined-inclusive intervals. However, we do not know a dichotomy classification for the instances of the problem that admit a PTAS. We ask for a dichotomy of the following form. If H belongs to class X of bipartite graphs, then there is a PTAS for Max-Min allocation problem. Otherwise, there is no PTAS.

Bibliography

- [1] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. *Journal of Scheduling*, 1:55–66, 1998.
- [2] C. Annamalai, C. Kalaitzis, and O. Svensson. Combinatorial algorithm for restricted max-min fair allocation. In *SODA*, 2015.
- [3] A. Asadpour, U. Feige, and A. Saberi. Santa claus meets hypergraph matchings. In *APPROX-RANDOM*, pages 10–20, 2008.
- [4] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. In *STOC*, pages 114–121. ACM, 2007.
- [5] N. Bansal and M. Sviridenko. The santa claus problem. In *STOC*, pages 31–40. ACM, 2006.
- [6] M. Bateni, M. Charikar, and V. Guruswami. Maxmin allocation via degree lower-bounded arborescences. In *STOC*. ACM, 2009.
- [7] I. Bezáková and V. Dani. Allocating indivisible goods. *SIGecom Exchanges*, 5(3):11–18, 2005.
- [8] K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using pq-tree algorithms. *J. Comput. Syst. Sci.*, 13(3):335–379, 1976.
- [9] A. Brandstädt, V. B. Le, and J. P. Spinrad. *Graph classes: a survey*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA, 1999.
- [10] D. Chakrabarty, J. Chuzhoy, and S. Khanna. On allocating goods to maximize fairness. In *FOCS*, pages 107–116, 2009.
- [11] T. Ebenlendr, M. Krcál, and J. Sgall. Graph balancing: A special case of scheduling unrelated parallel machines. *Algorithmica*, 68(1):62–80, 2014.
- [12] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [13] F. Glover. Maximum matching in a convex bipartite graph. *Naval research Logistics Quarterly*, 14:313–316, 1967.
- [14] G. Gutin, P. Hell, A. Rafiey, and A. Yeo. A dichotomy for minimum cost graph homomorphisms. *Eur. J. Comb.*, 29(4):900–911, 2008.
- [15] M. Habib, R. M. McConnell, C. Paul, and L. Viennot. Lex-bfs and partition refinement, with applications to transitive orientation, interval graph recognition and consecutive ones testing. *Theor. Comput. Sci.*, 234(1-2):59–84, 2000.
- [16] B. Haeupler, B. Saha, and A. Srinivasan. New constructive aspects of the lovász local lemma. *J. ACM*, 58(6):28, 2011.

- [17] P. R. Halmos and H. E. Vaughan. The marriage problem. *American Journal of Mathematics*, pages 214–215, 1950.
- [18] K. Khodamoradi, R. Krishnamurti, A. Rafiey, and G. Stamoulis. PTAS for ordered instances of resource allocation problems. In A. Seth and N. K. Vishnoi, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2013, December 12–14, 2013, Guwahati, India*, volume 24 of *LIPIcs*, pages 461–473. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [19] S. Khot and A. K. Ponnuswami. Approximation algorithms for the max-min allocation problem. In *APPROX-RANDOM*, pages 204–217, 2007.
- [20] K. Lee, J. Y. Leung, and M. L. Pinedo. A note on graph balancing problems with restrictions. *Inf. Process. Lett.*, 110(1):24–29, 2009.
- [21] J. K. Lenstra, D. B. Shmoys, and É. Tardos. Approximation algorithms for scheduling unrelated parallel machines. In *FOCS*, pages 217–224. IEEE, 1987.
- [22] M. Mastrolilli and G. Stamoulis. Restricted max-min fair allocations with inclusion-free intervals. In J. Gudmundsson, J. Mestre, and T. Viglas, editors, *Computing and Combinatorics - 18th Annual International Conference, COCOON 2012, Sydney, Australia, August 20–22, 2012. Proceedings*, volume 7434 of *Lecture Notes in Computer Science*, pages 98–108. Springer, 2012.
- [23] J. Meidanis, O. Porto, and G. P. Telles. On the consecutive ones property. *Discrete Applied Mathematics*, 88(1-3):325–354, 1998.
- [24] G. Muratore, U. M. Schwarz, and G. J. Woeginger. Parallel machine scheduling with nested job assignment restrictions. *Oper. Res. Lett.*, 38(1):47–50, 2010.
- [25] L. Polacek and O. Svensson. Quasi-polynomial local search for restricted max-min fair allocation. In *ICALP*, 2012.
- [26] B. Saha and A. Srinivasan. A new approximation technique for resource-allocation problems. In *ICS*, pages 342–357. Tsinghua University Press, 2010.
- [27] U. M. Schwarz. A ptas for scheduling with tree assignment restrictions. *CoRR*, abs/1009.4529, 2010.
- [28] J. Sgall. Randomized on-line scheduling of parallel jobs. *J. Algorithms*, 21(1):149–175, 1996.
- [29] J. Spinrad, A. Brandstädt, and L. Stewart. Bipartite permutation graphs. *Discrete Applied Mathematics*, 18(3):279 – 292, 1987.
- [30] J. P. Spinrad. *Efficient graph representations*. Fields Institute monographs. American Mathematical Society, Providence, RI, 2003.
- [31] O. Svensson. Santa claus schedules jobs on unrelated machines. In *STOC*, pages 617–626, 2011.
- [32] J. Verschae and A. Wiese. On the configuration-lp for scheduling on unrelated machines. *J. Scheduling*, 17(4):371–383, 2014.
- [33] G. Woeginger. A polynomial-time approximation scheme for maximizing the minimum machine completion time. *Operation Research Letters*, 20(4):149–154, 1997.